arXiv:1405.2294v1 [cs.LG] 25 Apr 2014

Nonparametric Detection of Anomalous Data via Kernel Mean Embedding [1]

Shaofeng Zou, Yingbin Liang, H. Vincent Poor, Xinghua Shi [2]

Abstract

An anomaly detection problem is investigated, in which there are totally n sequences with s anomalous sequences to be detected. Each normal sequence contains m independent and identically distributed (i.i.d.) samples drawn from a distribution p, whereas each anomalous sequence contains m i.i.d. samples drawn from a distribution q that is distinct from p. The distributions p and q are assumed to be unknown a priori. Two scenarios, respectively with and without a reference sequence generated by p, are studied. Distribution-free tests are constructed using maximum mean discrepancy (MMD) as the metric, which is based on mean embeddings of distributions into a reproducing kernel Hilbert space (RKHS). For both scenarios, it is shown that as the number n of sequences goes to infinity, if the value of s is known, then the number m of samples in each sequence should be at the order $\mathcal{O}(\log n)$ or larger in order for the developed tests to consistently detect s anomalous sequences. If the value of s is unknown, then m should be at the order strictly larger than $\mathcal{O}(\log n)$. Computational complexity of all developed tests is shown to be polynomial. Numerical results demonstrate that our tests outperform (or perform as well as) the tests based on other competitive traditional statistical approaches and kernel-based approaches under various cases. Consistency of the proposed test is also demonstrated on a real data set.

Key words: Anomaly detection, consistency, distribution-free tests, kernel mean embedding, maximum mean discrepancy (MMD), nonparametric model, reproducing kernel Hilbert space (RKHS).

1 Introduction

In this paper, we study an anomaly detection problem (see Figure 1), in which there are totally n sequences with s anomalous sequences to be detected. Each normal sequence consists of m independent and identically distributed (i.i.d.) samples drawn from a distribution

[1]The work of S. Zou and Y. Liang was supported by a National Science Foundation CAREER Award under Grant CCF-10-26565. The work of H. V. Poor was supported by National Science Foundation under Grants DMS-11-18605 and ECCS-13-43210.

[2]Shaofeng Zou and Yingbin Liang are with the Department of Electrical Engineering and Computer Science, Syracuse University, Syracuse, NY 13244 USA (email: {szou02,yliang06}@syr.edu). H. Vincent Poor is with the Department of Electrical Engineering, Princeton University, Princeton, NJ 08544 USA (email: poor@princeton.edu). Xinghua Shi is with the Department of Bioinformatics and Genomics, University of North Carolina at Charlotte, Charlotte, NC 28223 (email: xshi3@uncc.edu).

p, whereas each anomalous sequence contains i.i.d. samples drawn from a distribution q that is distinct from p. The distributions p and q are assumed to be unknown a priori. The goal is to build distribution-free tests to detect the s anomalous data sequences generated by q out of all data sequences.

The study of such a problem is very useful in many applications. For example, in cognitive wireless networks, signals follow different distributions either p or q depending on whether the channel is busy or vacant. A major issue in such a network is to identify vacant channels out of a large number of busy channels based on their corresponding signals in order to utilize vacant channels for improving spectral efficiency. This problem was studied in [1, 2] under the assumption that p and q are known, whereas in this paper, we study the problem with unknown p and q. Other applications include detecting anomalous DNA sequences out of normal sequences, detecting virus infected computers from other virus free computers, and detecting slightly modified images from other untouched images.

We note that in the model considered here, each data point contains a sequence of samples drawn from one distribution. This is different from the typical anomaly or outlier detection problems studied in machine learning [3, 4], in which each data point contains only one sample. The parametric model of the problem has been well studied, e.g., [1], in which it is assumed that the distributions of p and q are known a priori and can be exploited for detection. However, the nonparametric model is less explored, in which it is assumed that the distributions p and q are unknown and can be arbitrary. Recently, Li, Nitinawarat and Veeravalli proposed the divergence-based generalized likelihood tests in [5], and characterized the error decay exponents of these tests. However, [5] studied only the case when the distributions p and q are discrete with finite alphabets, and their tests utilize empirical probability mass functions of p and q.

In this paper, we study the nonparametric model, in which distributions p and q can be continuous and arbitrary. The major challenges to solve this problem (compared to the discrete case studied in [5]) lie in: (1) it is difficult to accurately estimate continuous distributions with limited samples for further anomaly detection; (2) it is difficult to design low complexity tests with continuous distributions; and (3) building distribution-free consistent tests (and further guaranteeing exponential error decay) is challenging for arbitrary distributions.

A number of statistical approaches and tools may be applied to solving this problem. A natural approach, e.g., FR-Smirnov test [6], is to first estimate the distributions based on data samples, and then compare the estimated distributions for anomaly detection. Such an approach typically does not perform very well, because it requires a large number of samples to accurately estimate continuous distributions, particularly when p and q represent distributions of vector random variables. Otherwise, the error in estimating the distributions can propagate to the anomaly detection step. Some traditional statistical approaches such as t-test, FR-Wolf test [6], and Hall test [7] do not require distribution estimation as an intermediate step, and can be applied to address this problem. However, t-test and FR-wolf test do not perform well for arbitrary distributions. Hall test has a large computational complexity. More recently, the kernel-based approaches such as kernel density ratio (KDR)

test [8] and kernel Fisher discriminant analysis (KFDA) test [9] have been developed, which use kernels to estimate certain distance metrics between two distributions. In particular, KDR test uses kernels to estimate the ratio between two probability densities and then further estimates the divergence between the two probability distributions.

In this paper, our approach that is introduced below falls into the class of kernel-based approaches. For the sake of comparison, we also develop/implement tests based on the traditional statistical approaches and other kernel-based approaches mentioned above, although those approaches were not applied specifically to our anomaly detection problem previously. We demonstrate that our tests outperform or equal those other tests under various testing cases.

More specifically, our approach adopts the emerging technique based on mean embedding of distributions into a reproducing kernel Hilbert space (RKHS) [10, 11]. The idea is to map probability distributions into a RKHS with an associated kernel such that distinguishing the two probabilities can be carried out by distinguishing their corresponding embeddings in the RKHS. Such an approach is justified by the fact that mapping of distributions into RKHS is injective (i.e., one-to-one) for certain kernels including Gaussian and Laplace kernels as shown in [11–14]. Since RKHS naturally carries a distance metric, mean embeddings of distributions can be compared easily based on their distances in the RKHS using the samples of distributions. Such a metric is referred to as the *maximum mean discrepancy (MMD)* as introduced in [15, 16]. MMD-based approaches have several advantages. For example, MMD can be easily estimated based on samples, and hence leads to low complexity tests. It is particularly advantageous for vector distributions. Furthermore, MMD-based approaches do not require estimation of probability density functions as an intermediate step, but directly estimate the distance of distributions to build tests, thus avoiding error propagation.

In this paper, we apply MMD as a metric to construct our tests for detecting data sequences generated by the anomalous distribution. We are interested in the large data regime, in which the total number n of data sequences goes to infinity. This is motivated by many practical applications such as DNA sequence detection, in which anomalous DNA sequences are required to be detected out of a large number of normal data sequences. It is clear that as the total number n of sequences becomes large, (and possibly the number s of anomalous data sequences also becomes large), it becomes increasingly challenging to consistently detect all anomalous sequences. It then requires that the number m of samples in each data sequence correspondingly enlarges in order to more accurately detect anomalous sequences. Hence, there is a tension between (n, s) and m in the asymptotic regime as n goes to infinity to guarantee correct detection. This also differentiates this paper from the study in [5], in which n and s are assumed to be fixed and only the number of samples becomes large.

We summarize our main contributions as follows.

(1) We construct MMD-based distribution-free consistent tests, which enjoy low computational complexity (which is polynomial as n increases) and superb performance (as compared to other tests described in (5)).

(2) We take a new perspective to analyze the consistency of tests. In contrast to consistency

analysis in classical statistical theory as in [5], which assumes that the problem dimension (i.e., the number n of sequences and the number s of anomalous sequences in our problem) is fixed and the sample size m increases, our focus is on the regime in which the problem dimension increases. We are interested in characterizing how the sample size m should scale with n and s in order to guarantee the consistency of our tests. We show that in general the sample size m can be much smaller than n (i.e., in the order of $\log n$), which is desirable.

(3) We perform a comprehensive study of the problem in various cases. We study two scenarios, with and without a reference data sequence generated by p, respectively. For each scenario, we build distribution-free tests for two cases with and without the knowledge of the number s of anomalous sequences. Comparing the performance of these tests suggests that a reference sequence does not always provide advantage for reducing the number m of samples needed for consistent detection, but in general helps to reduce computational complexity. Nevertheless, lack of the information about the number s of anomalous sequences results in order-level increase in m needed for consistent detection. For the scenario without a reference sequence, we further develop low complexity consistent tests by exploiting the asymptotic behavior of s and n.

(4) We provide an interesting example study, in which the distribution q is the mixture of the distribution p and the anomalous distribution $\tilde{q}$. In such a case, the anomalous sequence contains only sparse samples from the anomalous distribution. Our results for such a model quantitatively characterize the impact of the sparsity level of anomalous samples on the scaling behavior of the sample size m, in order to guarantee consistency of the proposed tests.

(5) We provide extensive numerical results. We first provide numerical results to demonstrate our theoretical assertions. We then develop/implement tests based on a number of competitive traditional statistical approaches and kernel-based approaches, and compare our MMD-based tests with these tests. We further compare our tests with the divergence-based generalized likelihood test proposed in [5]. Our numerical results demonstrate that the MMD-based test is the best performed test (or among the best performed tests) under various experimental cases. We finally apply our test to a real data set to demonstrate that our test is consistent and among the best performed tests.

The rest of the paper is organized as follows. In Sections 2, we describe the problem formulation and introduce the MMD-based approach. In Sections 3 and 4, we present our tests and theoretical results on the performance guarantee of these tests respectively for the cases with and without a reference sequence. In Section 5, we provide numerical results. Finally in Section 6, we conclude our paper with remarks on future work.

2 Problem Statement and Preliminaries on MMD

In this section, we first describe the problem that we study, and then briefly introduce the idea of kernel mean embedding of distributions and the metric of MMD that we use for building our tests.

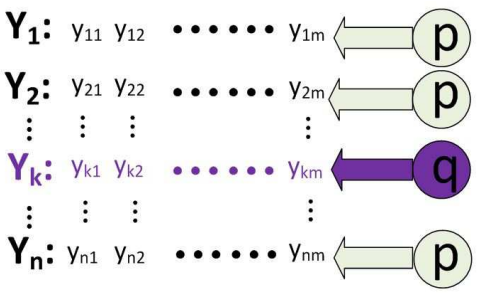

Figure 1: An anomaly detection model with data sequences generated by distribution p and anomalous distribution q.

2.1 Problem Statement

We study an anomaly detection problem (see Figure 1), in which there are in total n data sequences denoted by Y_k for $1 \leq k \leq n$. Each data sequence Y_k consists of m i.i.d. samples $y_{k1}, \ldots, y_{km}$ drawn from either a distribution p or an anomalous distribution q, where $p \neq q$. In the sequel, we use the notation $Y_k := (y_{k1}, \ldots, y_{km})$. We assume that the distributions p and q are arbitrary and unknown a priori. Our goal is to build distribution-free tests to detect data sequences generated by the anomalous distribution q.

We assume that s out of n data sequences are anomalous, i.e., are generated by the anomalous distribution q. We study both cases with the value of s known and unknown a priori, respectively. We are interested in the asymptotical regime, in which the number n of data sequences goes to infinity. We assume that the number s of anomalous sequences satisfies $\frac{s}{n} \to \alpha$ as $n \to \infty$, where $0 \leq \alpha \leq 1$. This includes the following three cases: (1) s is fixed as $n \to \infty$; (2) $s \to \infty$, but $\frac{s}{n} \to 0$ as $n \to \infty$; and (3) $\frac{s}{n}$ approaches to a positive constant, which is less than or equal to 1. Some of our results are applicable to the case with $s = 0$, i.e., the null hypothesis in which there is no anomalous sequence. We will comment on such a case when the corresponding results are presented. In this paper, $f(n) = \mathcal{O}(g(n))$ denotes $f(n)/g(n)$ converges to a constant as $n \to \infty$.

We study two scenarios. Scenario 1 assumes that there is a reference data sequence X available a priori, which contains i.i.d. samples $(x_1, \ldots, x_m)$ generated from the distribution p. In the sequel, we use the notation $X := (x_1, \ldots, x_m)$. Scenario 2 assumes that such a reference data sequence X is not available. Scenario 1 is reasonable because it is often the case in applications that samples from distribution p are easy to collect, and can hence be exploited to detect anomalous data sequences.

We next define the probability of detection error as the performance measure of tests. We let $\mathcal{I}$ denote the set that contains indices of all anomalous data sequences. Hence, the cardinality $|\mathcal{I}| = s$. We let $\hat{\mathcal{I}}^n$ denote a sequence of index sets that contain indices of all anomalous data sequences claimed by a corresponding sequence of tests.

Definition 1. *A sequence of tests are said to be consistent if*

$$\lim_{n\to\infty} P_e = \lim_{n\to\infty} P\{\hat{\mathcal{I}}^n \neq \mathcal{I}\} = 0. \tag{1}$$

We note that the above definition of consistency is with respect to the number n of sequences instead of the number m of samples. However, as n becomes large (and possibly as s becomes large), it is increasingly challenging to consistently detect all anomalous data sequences. It then requires that the number m of samples becomes large enough in order to more accurately detect anomalous sequences. Therefore, the limit in the above definition in fact refers to the asymptotic regime, in which m scales fast enough as n goes to infinity in order to guarantee asymptotically small probability of error.

Furthermore, for a consistent test, it is also desirable that the error probability decays exponentially fast with respect to the number m of samples.

Definition 2. *A sequence of tests are said to be exponentially consistent if*

$$\liminf_{m\to\infty} -\frac{1}{m}\log P_e = \liminf_{m\to\infty} -\frac{1}{m}\log P\{\hat{\mathcal{I}}^n \neq \mathcal{I}\} > 0. \tag{2}$$

The above definition is with respect to the asymptotic regime of m, because the exponential decay behavior of the error probability is with respect to m. In fact, the limit refers to the asymptotic regime, in which m scales fast enough as n goes to infinity in order to guarantee exponential error decay.

In this paper, our goal is to construct distribution-free tests for detecting anomalous sequences, and characterize the scaling behavior of m with n (and possibly s) so that the developed tests are consistent (and possibly exponentially consistent).

An example with sparse anomalous samples. In this paper, we also study an interesting example, in which the distribution q is a mixture of the distribution p with probability $1-\epsilon$ and an anomalous distribution $\tilde{q}$ with probability ϵ, where $0 < \epsilon \leq 1$, i.e., $q = (1-\epsilon)p+\epsilon\tilde{q}$. It can be seen that if ϵ is small, the majority of samples in an anomalous sequence are drawn from the distribution p, and only sparse samples are drawn from the anomalous distribution q. The value of ϵ captures the sparsity level of anomalous samples. Here, ϵ can scale as n increases, and is hence denoted by ϵ_n. We study how ϵ_n affects the number of samples needed for consistent detection.

2.2 Introduction of MMD

In this subsection, we briefly introduce the idea of mean embedding of distributions into RKHS [10,11] and the metric of MMD. Suppose $\mathcal{P}$ is a set that includes a class of probability distributions, and suppose $\mathcal{H}$ is the RKHS with an associated kernel $k(\cdot,\cdot)$ (see [17, 18] for an introduction of RKHS). We define a mapping from $\mathcal{P}$ to $\mathcal{H}$ such that each distribution $p \in \mathcal{P}$ is mapped to an element in $\mathcal{H}$ as follows

$$\mu_p(\cdot) = \mathbb{E}_p[k(\cdot, x)] = \int k(\cdot, x)dp(x).$$

Here, $\mu_p(\cdot)$ is referred to as the *mean embedding* of the distribution p into the Hilbert space $\mathcal{H}$. Due to the reproducing property of $\mathcal{H}$, it is clear that $\mathbb{E}_p[f] = \langle \mu_p, f \rangle_{\mathcal{H}}$ for all $f \in \mathcal{H}$.

It is desirable that the embedding is *injective* such that each $p \in \mathcal{P}$ is mapped to a unique element $\mu_p \in \mathcal{H}$. It has been shown in [11–14] that for many RKHSs such as those associated with Gaussian and Laplace kernels, the mean embedding is injective. In order to distinguish between two distributions p and q, [16] introduced the following quantity of maximum mean discrepancy (MMD) based on the mean embeddings μ_p and μ_q of p and q in RKHS:

$$\mathrm{MMD}[p, q] := \|\mu_p - \mu_q\|_{\mathcal{H}}. \tag{3}$$

It is clear that the MMD equals the distance between the mean embeddings μ_p and μ_q of the two distributions p and q, and can be used to distinguish between p and q due to the injectivity of mean embedding.

It can also be shown that

$$\mathrm{MMD}[p, q] = \sup_{\|f\|_{\mathcal{H}} \leq 1} \mathbb{E}_p[f(x)] - \mathbb{E}_q[f(x)].$$

Namely, $\mathrm{MMD}[p, q]$ achieves the maximum of the mean difference of a function between the two distributions over all unit-norm functions in the RKHS $\mathcal{H}$.

Due to the reproducing property of kernel, it can be easily shown that

$$\mathrm{MMD}^2[p, q] = \mathbb{E}_{x,x'}[k(x, x')] - 2\mathbb{E}_{x,y}[k(x, y)] + \mathbb{E}_{y,y'}[k(y, y')], \tag{4}$$

where x and x' have independent but the same distribution p, and y and y' have independent but the same distribution q. An unbiased estimator of $\mathrm{MMD}^2[p, q]$ based on l_1 samples of X and l_2 samples of Y is given as follows,

$$\mathrm{MMD}_u^2[X, Y] = \frac{1}{l_1(l_1 - 1)} \sum_{i=1}^{l_1} \sum_{j \neq i}^{l_1} k(x_i, x_j) + \frac{1}{l_2(l_2 - 1)} \sum_{i=1}^{l_2} \sum_{j \neq i}^{l_2} k(y_i, y_j) - \frac{2}{l_1 l_2} \sum_{i=1}^{l_1} \sum_{j=1}^{l_2} k(x_i, y_j). \tag{5}$$

We note that other estimators of the $\mathrm{MMD}^2[p, q]$ are also available, which can be used for our problem. In this paper, we focus on the unbiased estimator given above to convey the central idea.

3 Scenario I: with Reference Sequence

In this section, we study the scenario, in which a reference sequence X containing samples generated from the distribution p is available. We start with a simple case, in which there is only one anomalous sequence out of the total n sequences, i.e. $s = 1$. We then generalize our study to the case, in which $1 \leq s \leq n-1$, and s is known a priori. We further study the case with the value of s unknown. We finally provide an example, in which anomalous sequences

contain only sparse samples generated by an anomalous distribution. We construct tests for all cases and analyze the performance and complexity of these tests. In particular, we characterize how m scales with n and s in order to guarantee consistent detection of all anomalous data sequences.

We first consider the case with only one anomalous sequence. With the reference sequence X, we compute $\mathrm{MMD}_u^2[X, Y_k]$ for each sequence Y_k for $1 \leq k \leq n$ using (5). It is clear that if Y_k is generated by p, then $\mathrm{MMD}_u^2[X, Y_k]$ is an estimator of $\mathrm{MMD}^2[p, p]$ and hence should be close to zero. If Y_k is anomalous, $\mathrm{MMD}_u^2[X, Y_k]$ is an estimator of $\mathrm{MMD}^2[p, q]$, which is a positive constant. This understanding naturally leads to the following distribution-free test when $s = 1$. The sequence k^* is the index of the anomalous data sequence if

$$k^* = \arg \max_{1 \leq k \leq n} \mathrm{MMD}_u^2[X, Y_k]. \tag{6}$$

The following theorem characterizes the condition under which the above test is consistent.

Theorem 1. *Consider the anomaly detection model with a reference sequence generated by p and with one anomalous sequence. Suppose the test* (6) *applies a bounded kernel with $0 \leq k(x, y) \leq K$ for any (x, y). Then the test* (6) *is consistent if*

$$m > \frac{24K^2(1+\eta)}{MMD^4[p, q]} \log n, \tag{7}$$

where η is any positive constant. Furthermore, under the above condition, the test (6) *is also exponentially consistent.*

Proof. The proof of Theorem 1 is given in Appendix A. □

We note that the boundedness condition $0 \leq k(x, y) \leq K$ on the kernel function is satisfied by many kernels such as Gaussian kernel and Laplace kernel. Theorem 1 implies that it is sufficient to have $\mathcal{O}(\log n)$ samples in each data sequence in order to guarantee consistency of the test (6). This is desirable in practice because as the number of sequences gets large (and hence detection becomes more challenging), the number of sample needed for building a distribution-free and consistent test can still be much smaller, i.e., in the logarithmic order of the number of sequences.

We now consider the general case with s anomalous sequences, where $1 \leq s \leq n-1$. Here, we allow $s \geq \frac{n}{2}$ for generality of our result. We first consider the case when the value of s is known a priori, and build the following test. We compute $\mathrm{MMD}_u^2[X, Y_k]$ for each $1 \leq k \leq n$, and pick sequences with the largest s values of $\mathrm{MMD}_u^2[X, Y_k]$ to be anomalous. More specifically, the test outputs the following set that contains indices of anomalous sequences:

$$\begin{aligned} \hat{\mathcal{I}} = &\{k : \mathrm{MMD}_u^2[X, Y_k] \text{ is among the } s \text{ largest} \\ &\text{values of } \mathrm{MMD}_u^2[X, Y_i] \text{ for } i = 1, \ldots, n\}. \end{aligned} \tag{8}$$

The following theorem characterizes the condition under which the above test is consistent.

Theorem 2. *Consider the anomaly detection model with a reference sequence generated by p and with s anomalous sequence, where $1 \leq s \leq n-1$. Assume the value of s is known. Further assume that the test* (8) *applies a bounded kernel with $0 \leq k(x,y) \leq K$ for any (x,y). Then the distribution-free test* (8) *is consistent if*

$$m > \frac{24K^2(1+\eta)}{MMD^4[p,q]} \log((n-s)s), \tag{9}$$

where η is any positive constant. Furthermore, under the above condition, the test (8) *is also exponentially consistent. The computational complexity of the test* (8) *is* $\max\{\mathcal{O}(nm^2), \mathcal{O}(ns - \frac{s^2}{2})\}$.

Proof. See Appendix B. □

Since s has an order less than or equal to $\mathcal{O}(n)$, Theorem 2 implies that it is sufficient to have $\mathcal{O}(\log n)$ samples in each data sequence in order to guarantee consistency of the test. Although s does not affect the order of m, it is still interesting to further understand how s affects the exact value of the threshold in (9). It can be seen that if $s \leq \frac{n}{2}$, the threshold on m to guarantee consistent detection increases as s increases, which is reasonable. It is somewhat surprising that if $s > \frac{n}{2}$, the threshold on m decreases as s increases. This is in fact also intuitive, because if $s > \frac{n}{2}$, the number of normal sequences is less than $\frac{n}{2}$, and it is hence more convenient to detect normal sequences (and consequently anomalous sequences are also identified). As s increases, the number of normal sequences decreases, and thus the detection is easier.

Theorem 2 also implies that computational complexity of the test (8) is at most $\mathcal{O}(n^2)$, because the order $\mathcal{O}(\log n)$ for m is sufficient for the test to be consistent, and s is at most n.

We next consider the case, in which the value of s is unknown. For this case, the test (8) is not applicable, because it depends on the value of s. In order to build a test now, we observe that for large enough m, $\mathrm{MMD}_u^2[X, Y_k]$ should be close to zero if Y_k is drawn from p, and should be far away enough from zero (in fact, close to $\mathrm{MMD}^2[p,q]$) if Y_k is drawn from the anomalous distribution q. Based on this understanding, we build the following test:

$$\widehat{\mathcal{I}} = \{k : \mathrm{MMD}_u^2[X, Y_k] > \delta_n\} \tag{10}$$

where the threshold $\delta_n \to 0$ as $n \to \infty$. This requirement on δ_n is to guarantee that the threshold is asymptotically less than $\mathrm{MMD}^2[p,q]$, which is positive but unknown. The following theorem characterizes the condition that m should satisfy in order for the test (10) to be consistent.

Theorem 3. *Consider the anomaly detection model with a reference sequence generated by p and with s anomalous sequences, where $0 \leq s \leq n-1$. Assume that the value of s is unknown a priori. Further assume that the test* (10) *adopts a threshold δ_n that converges to*

0 *as* $n \to \infty$*, and applies a bounded kernel with* $0 \leq k(x, y) \leq K$ *for any* (x, y)*. Then the test* (10) *is consistent if*

$$m > 16(1+\eta)K^2 \max\left\{\frac{\log(\max\{s,1\})}{(MMD^2[p,q]-\delta_n)^2}, \frac{\log(n-s)}{\delta_n^2}\right\} \tag{11}$$

where η *is any positive constant. The computational complexity of the test* (10) *is* $\mathcal{O}(nm^2)$*.*

Proof. See Appendix C. □

Remark 1. *Theorem 3 is also applicable to the case with* $s = 0$*, i.e., the null hypothesis when there is no anomalous sequence.*

Theorem 3 implies that if $\frac{s}{n} < 1$ as $n \to \infty$, the threshold on m in (11) is dominated by the second term. Since $\delta_n \to 0$ as $n \to \infty$, m should scale strictly faster than $\mathcal{O}(\log n)$ in order to guarantee consistent detection. Compared to the case with the value of s known (for which it is sufficient for m to scale at the order $\mathcal{O}(\log n)$), the threshold on m has order-level increase due to lack of the knowledge of s. An extreme case occurs if $\frac{s}{n} = 1$ as $n \to \infty$, in which the order $\mathcal{O}(\log n)$ for m is sufficient. This is reasonable, because now anomalous sequences dominate so that errors caused by the asymptotically small threshold δ_n do not dominate the performance, and hence do not enlarge the requirement on the order of m.

We also note that the test (10) is not exponentially consistent. In fact, when there is no null hypothesis (i.e., $s > 1$), an exponentially consistent test can be built as follows. For each subset $\mathcal{S}$ of $\{1, \ldots, n\}$, we compute the average $\frac{1}{|\mathcal{S}|}\sum_{k\in\mathcal{S}} \mathrm{MMD}_u^2[X, Y_k]$, and the test finds the set of indices corresponding to the largest average value. However, for such a test to be consistent, m needs to scale linearly with n and computational complexity is exponential with n, which are not desirable.

We finally study the example with sparse anomalous samples that we introduce in Section 2.1. In this model, ϵ_n captures the sparsity level of anomalous samples. The following results capture the impact of ϵ_n on the scaling behavior of m to guarantee consistency of our test.

Corollary 1. *Consider the model with the distribution* p *and the anomalous distribution* $q = (1-\epsilon_n)p + \epsilon_n \tilde{q}$*, where* $0 < \epsilon_n \leq 1$*. If* s *is known, then the test* (8) *is consistent if*

$$m > \frac{24K^2(1+\eta)}{MMD^4[p,\tilde{q}]\epsilon_n^4}\log(s(n-s)), \tag{12}$$

where η *is any positive constant.*

If s *is unknown, then the test* (10) *is consistent if*

$$m > 16(1+\eta)K^2 \max\left\{\frac{\log(\max\{s,1\})}{(\epsilon_n^2 MMD^2[p,\tilde{q}]-\delta_n)^2}, \frac{\log(n-s)}{\delta_n^2}\right\} \tag{13}$$

where η *is any positive constant, and* $\delta_n < \epsilon_n^2 MMD^2[p,\tilde{q}]$ *for large enough* n*.*

We note that for the case with s unknown, the design of the test (10) requires information about the scaling behavior of ϵ_n, which is usefully available in practice.

Proof. The proof follows from Theorems 2 and 3 by substituting:

$$\begin{aligned}
&\mathrm{MMD}^2[p,q] \\
&= \mathbb{E}_{x,x'}[k(x,x')] - 2\mathbb{E}_{x,y}[k(x,y)] + \mathbb{E}_{y,y'}[k(y,y')] \\
&= \mathbb{E}_{x,x'}[k(x,x')] - 2(1-\epsilon_n)\mathbb{E}_{x,x'}[k(x,x')] - 2\epsilon_n\mathbb{E}_{x,\tilde{y}}[k(x,\tilde{y})] + (1-\epsilon_n)^2\mathbb{E}_{x,x'}[k(x,x')] \\
&\quad + 2\epsilon_n(1-\epsilon_n)\mathbb{E}_{x,\tilde{y}}[k(x,\tilde{y})] + \epsilon_n^2\mathbb{E}_{\tilde{y},\tilde{y}'}[k(\tilde{y},\tilde{y}')] \\
&= \epsilon_n^2\mathrm{MMD}^2[p,\tilde{q}], \qquad (14)
\end{aligned}$$

where x and x' are independent but have the same distribution p, y and y' are independent but have the same distribution q, and $\tilde{y}$ and $\tilde{y}'$ are independent but have the same distribution $\tilde{q}$. □

Corollary 1 implies that if ϵ_n is a constant, then the scaling behavior of m needed for consistent detection does not change. However, if $\epsilon_n \to 0$ as $n \to \infty$, i.e., anomalous sequences contain more sparse anomalous samples, then m needs to scale faster with n in order to guarantee consistent detection. This is reasonable because the number m of samples in each sequence should have a higher order to cancel out the impact of the increasingly sparse anomalous samples in each anomalous sequence. Corollary 1 explicitly captures such tradeoff between the number m of samples and the sparsity level ϵ_n of anomalous samples in addition to the number n of total data sequences and the number s of anomalous sequences.

4 Scenario II: Without Reference Sequence

In this section, we consider the case, in which a reference sequence generated by the distribution p is not available. As for the scenario with a reference sequence, here we start with a simple case with $s = 1$, and then study the more general case, in which $\frac{s}{n} \to \alpha$ as $n \to \infty$, where $0 \leq \alpha \leq 1$. We study both cases with the value of s known and unknown, respectively. We finally study the example, with sparse anomalous samples.

We first study the case with $s = 1$. For each sequence Y_k, we use $\overline{Y}_k$ to denote the $(n-1)m$ dimensional sequence that stacks all other sequences together, as given by

$$\overline{Y}_k = \{Y_1, \ldots, Y_{k-1}, Y_{k+1}, \ldots, Y_n\}.$$

We then compute $\mathrm{MMD}_u^2[Y_k, \overline{Y}_k]$ for $1 \leq k \leq n$. It is clear that if Y_k is an anomalous sequence, then $\overline{Y}_k$ is fully composed of normal sequences. Hence, $\mathrm{MMD}_u^2[Y_k, \overline{Y}_k]$ is a good estimator of $\mathrm{MMD}^2[p,q]$, which is a positive constant. Nonetheless, if Y_k is a normal sequence, $\overline{Y}_k$ is composed of $n-2$ sequences generated by p and only one sequence generated by q. As n increases, the impact of the anomalous sequence on $\overline{Y}_k$ is negligible, and $\mathrm{MMD}_u^2[Y_k, \overline{Y}_k]$

should be close to zero. Based on the above understanding, we construct the following test when $s = 1$. The sequence k^* is the index of the anomalous data sequence if

$$k^* = \arg \max_{1 \le k \le n} \mathrm{MMD}_u^2[Y_k, \overline{Y}_k]. \tag{15}$$

The following theorem characterizes the condition under which the above test is consistent.

Theorem 4. *Consider the anomaly detection model without a reference sequence and with one anomalous sequence. Suppose the test* (15) *applies a bounded kernel with* $0 \le k(x, y) \le K$ *for any* (x, y)*. Then, the test* (15) *is consistent if*

$$m > \frac{16K^2(1+\eta)}{MMD^4[p,q]} \log n, \tag{16}$$

where η *is any positive constant. Furthermore, under the above condition, the test* (15) *is also exponentially consistent.*

Proof. See Appendix D. □

Theorem 4 implies that for the scenario without a reference sequence, $\mathcal{O}(\log n)$ samples are still sufficient to guarantee consistent detection.

We now consider the more general case in which $\frac{s}{n} \to \alpha$ as $n \to \infty$, where $0 \le \alpha < \frac{1}{2}$. Although we focus on the case with $\alpha < \frac{1}{2}$, the case with $\alpha > \frac{1}{2}$ is similar, with the roles of p and q exchanged. We first study the case, in which the value of s is known. Our test is a natural generalization of the test (15) except now the test picks the sequences with the largest s values of $\mathrm{MMD}_u^2[Y_k, \overline{Y}_k]$, which is given by

$$\begin{aligned} \hat{\mathcal{I}} = &\{k : \mathrm{MMD}_u^2[Y_k, \overline{Y}_k] \text{ is among the } s \text{ largest} \\ &\text{values of } \mathrm{MMD}_u^2[Y_i, \overline{Y}_i] \text{ for } i = 1, \ldots, n\}. \end{aligned} \tag{17}$$

The following theorem characterizes the condition under which the above test is consistent.

Theorem 5. *Consider the anomaly detection model without a reference sequence and with* s *anomalous sequences, where* $\frac{s}{n} \to \alpha$ *as* $n \to \infty$ *and* $0 \le \alpha < \frac{1}{2}$*. Assume the value of* s *is known. Further assume that the test* (17) *applies a bounded kernel with* $0 \le k(x, y) \le K$ *for any* (x, y)*. Then the test* (17) *is consistent if*

$$m > \frac{16K^2(1+\eta)}{(1-2\alpha)^2 MMD^4[p,q]} \log(s(n-s)), \tag{18}$$

where η *is any positive constant. Furthermore, under the above condition, the test* (17) *is also exponentially consistent. The computational complexity of the test* (17) *is* $\mathcal{O}(n^3 m^2)$*.*

Proof. See Appendix E. □

Theorem 5 implies that without a reference sequence, the test (17) requires $\mathcal{O}(\log n)$ samples in each data sequence in order to guarantee consistency of the test, which is the same as the test (8) with a reference sequence.

We further note that Theorem 5 is also applicable to the case in which $\alpha > \frac{1}{2}$ simply with the roles of p and q exchanged.

Remark 2. *Comparison of Theorems 5 and 2 suggests that if α is small enough, then the test* (17) *(that does not exploit a reference sequence) requires a smaller number of samples for guaranteeing consistent detection than the test* (8) *that uses a reference sequence. Although this seems counter-intuitive, it is in fact reasonable, because the test* (17) *uses $MMD_u^2[Y_i, \overline{Y}_i]$ as an estimator of MMD^2. The sequence $\overline{Y}_i$ contains more samples than the reference sequence, and hence leads to a more accurate estimator of MMD^2 if anomalous sequences have only negligible impact on $\overline{Y}_i$ (which holds if α is small). As α becomes larger, the sequence $\overline{Y}_i$ is contaminated by more anomalous sequences, and hence the performance of the test* (17) *eventually gets worse than the test* (8)*.*

Theorem 5 also suggests that computational complexity of the test (17), although is still polynomial in n, is larger than that of the test (8) with a reference sequence. Hence, a reference sequence in general helps to reduce computational complexity. Furthermore, as we remark below, we can construct a test without a reference sequence that has lower complexity than the test (17) (although the improved complexity is still larger than the test (8) with a reference sequence).

Remark 3. *For the case with $\frac{s}{n} \to 0$, as $n \to \infty$, we can build a test without a reference sequence that has reduced computational complexity. For each Y_k, instead of using $n-1$ sequences to build $\overline{Y}_k$ as in the test* (17)*, we take any l sequences out of the remaining $n-1$ sequences to build a sequence $\widetilde{Y}_k$, such that $\frac{l}{n} \to 0$ and $\frac{s}{l} \to 0$ as $n \to \infty$. Such an l exists for any s and n satisfying $\frac{s}{n} \to 0$ (e.g., $l = \sqrt{sn}$). It can be shown that using $\widetilde{Y}_k$ to replace $\overline{Y}_k$ in the test* (17) *still leads to consistent detection under the same condition given in Theorem 5. The computational complexity of such a test becomes $\mathcal{O}(nl^2m^2)$, which is substantially smaller than $\mathcal{O}(n^3m^2)$ of the test* (17)*, considering that l is less than n in the order sense.*

We next consider the case, in which the value of s is unknown a priori. Based on the understanding similar to scenario 1, we build the following test:

$$\widehat{\mathcal{I}} = \{k : \mathrm{MMD}_u^2[Y_k, \overline{Y}_k] > \delta_n\} \tag{19}$$

where $\delta_n \to 0$ and $\frac{s^2}{n^2\delta_n} \to 0$ as $n \to \infty$. We note that the above requirements on δ_n implies that the test (19) is applicable only when $\frac{s}{n} \to 0$ as $n \to \infty$. This includes two cases: (1) s is fixed and (2) $s \to \infty$ and $\frac{s}{n} \to 0$ as $n \to \infty$. Furthermore, the scaling behavior of s as n increases needs to be known in order to pick δ_n for the test. This is reasonable to assume because mostly in practice the scale of anomalous data points can be estimated based on domain knowledge.

The following theorem characterizes the condition under which the test (19) is consistent.

Theorem 6. *Consider the anomaly detection model without a reference sequence and with s anomalous sequences, where $\frac{s}{n} \to 0$, as $n \to \infty$. We assume that the value of s is unknown a priori. Further assume that the test (19) adopts a threshold δ_n such that $\delta_n \to 0$ and $\frac{s^2}{n^2\delta_n} \to 0$, as $n \to \infty$, and the test applies a bounded kernel with $0 \leq k(x,y) \leq K$ for any (x,y). Then the test (19) is consistent if*

$$m > 16(1+\eta)K^2 \max\Big\{\frac{\log(\max\{1,s\})}{(MMD^2[p,q]-\delta_n)^2}\,, \frac{\log(n-s)}{(\delta_n - \mathbb{E}\big[MMD_u^2[Y,\overline{Y}]\big])^2}\Big\}, \tag{20}$$

where η is any positive constant, and $E[MMD_u^2[Y,\overline{Y}]]$ is a constant, where Y is a sequence generated by p and $\overline{Y}$ is a stack of $(n-1)$ sequences with s sequences generated by q and the remaining sequences generated by p. The computational complexity of the test (19) is $\mathcal{O}(n^3m^2)$.

Proof. See Appendix F. □

We note that Theorem 6 is also applicable to the case with $s = 0$, i.e., the null hypothesis when there is no anomalous sequence. We further note that the test (19) is not exponentially consistent. Following the idea similar to scenario 1 with a reference sequence, for the case with $s > 0$, an exponentially consistent test can be built. However, the consistency of such a test requires that m scales linearly with n and the computational complexity is exponential with respect to n, which are not desirable.

Theorem 6 implies that the threshold on m to guarantee consistent detection has an order strictly larger than $\mathcal{O}(\log n)$, because $\frac{s}{n} \to 0$ and $\delta_n \to 0$ as $n \to \infty$. As we comment for scenario 1, this is the price paid due to lack of the knowledge of the value of s.

As for scenario 1 with a reference sequence, we finally study the example with sparse anomalous samples, i.e., the anomalous distribution $q = (1-\epsilon_n)p + \epsilon_n\tilde{q}$. The following result characterizes the impact of sparsity level ϵ_n on the scaling behavior of m to guarantee consistent detection.

Corollary 2. *Consider the model with the distribution p and the anomalous distribution $q = (1-\epsilon_n)p + \epsilon_n\tilde{q}$, where $0 < \epsilon_n \leq 1$. If s is known, then the test in (17) is consistent if*

$$m > \frac{16K^2(1+\eta)}{(1-2\alpha)^2\epsilon_n^4 MMD^4[p,\tilde{q}]}\log(s(n-s)), \tag{21}$$

where η is any positive constant.

If s is unknown, then the test in (19) is consistent if

$$m > 16(1+\eta)K^2 \max\Big\{\frac{\log(\max\{1,s\})}{(\epsilon_n^2 MMD^2[p,\tilde{q}]-\delta_n)^2}\,, \frac{\log(n-s)}{(\delta_n - \mathbb{E}\big[MMD_u^2[Y,\overline{Y}]\big])^2}\Big\}, \tag{22}$$

where η is any positive constant, $\delta_n < \mathbb{E}\big[MMD_u^2[Y,\overline{Y}]\big]$ for large enough n, Y is a sequence generated by p, and $\overline{Y}$ is a stack of $(n-1)$ sequences with s sequences generated by $\tilde{q}$ and the remaining sequences generated by p.

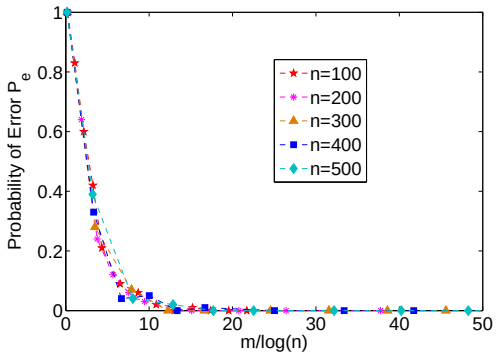

Figure 2: The performance of the test for the case with a reference sequence and with $s = 1$.

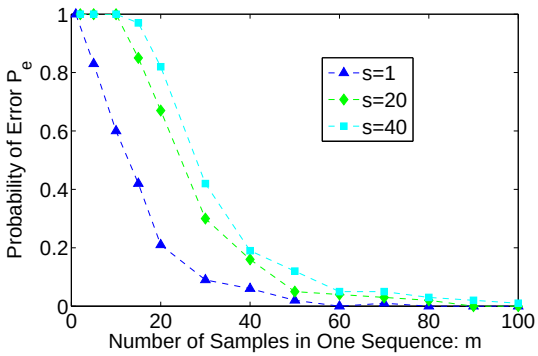

Figure 3: The performance of the test for the case with a reference sequence and with $s \geq 1$.

Proof. The proof follows from Theorems 5 and 6 by substituting $\mathrm{MMD}^2[p, q] = \epsilon_n^2 \mathrm{MMD}^2[p, \tilde{q}]$. □

5 Numerical Results

In this section, we provide numerical results to demonstrate our theoretical assertions, and compare our MMD-based tests with a number of tests based on other approaches. We also apply our test to a real data set.

5.1 Demonstration of Theorems

We start with scenario 1 with a reference sequence. We choose the distribution p to be Gaussian with mean zero and variance one, i.e., $\mathcal{N}(0, 1)$, and choose the anomalous distribution q to be Laplace distribution with mean one and variance one. We use the Gaussian kernel $k(x, x') = \exp(-\frac{|x-x'|^2}{2\sigma^2})$ with $\sigma = 1$.

For the case with $s = 1$, we run the test (6) for five cases with $n = 100, 200, 300, 400, 500$, respectively. In Figure 2, we plot how the probability of error changes with m. For illustrational convenience, we normalize m by $\log n$, i.e., the horizontal axis represents $\frac{m}{\log n}$. It is clear from the figure that when $\frac{m}{\log n}$ is above a certain threshold, the probability of error converges to zero, which is consistent with our theoretical results. Furthermore, for different values of n, all curves drop to zero almost at the same threshold. This observation confirms Theorem 1, which states that the threshold on $\frac{m}{\log n}$ depends only on the bound K of the kernel and MMD of the two distributions. Both quantities are constant for all values of n.

For the case with $s \geq 1$, we set $n = 100$, and run the test (8) for three cases with the numbers of anomalous sequences being $s = 1, 20, 40$, respectively. In Figure 3, we plot the probability of error as a function of m. It can be seen that for each value of s, when m is

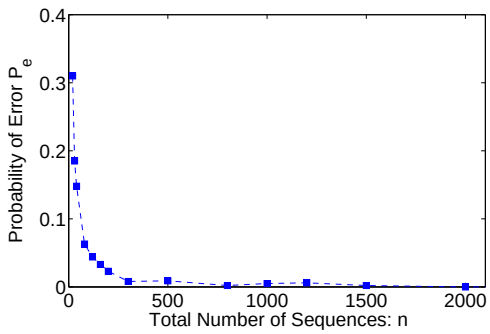

Figure 4: The performance of the test with a reference sequence and with the value of s unknown.

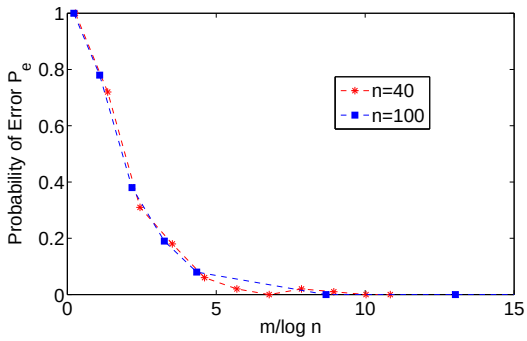

Figure 5: The performance of the test without a reference sequence and with $s = 1$.

above a certain threshold, the probability of error converges to zero, representing that our test is consistent. Furthermore, the threshold on m at which a curve drops to zero increases with s. This observation is consistent with Theorem 2, which suggests that the threshold on m increases with s if $s < \frac{n}{2}$.

We next study the case with a reference sequence and with s anomalous sequences, but here the value of s is unknown a priori. For this simulation, we set $s = 10$, but our test does not exploit such information. We choose the distribution p to be $\mathcal{N}(0, 0.5)$, and choose the distribution q to be a mixture of two Laplace distributions with equal probability: one with mean -3 and variance 0.5 and the other with mean 3 and variance 0.5. We apply the test (10), and set the threshold $\delta_n = \frac{1}{(\log n)^{0.7}}$, which converges to zero as $n \to \infty$. We set $m = \lceil 0.28(\log n)^{1.4} \log(n - s) \rceil$. In Figure 4, we plot the probability of error as a function of n. It can be seen that as n increases, the probability of error converges to zero. This clearly confirms Theorem 3, which implies that the chosen scaling behavior of m should guarantee consistency of the test.

We further study scenario 2, in which there is no reference sequence. We choose the distributions p and q to be the same as those for Figure 2, and use the same Gaussian kernel for our test here so that it is easy to compare the performance with scenario 1. We choose $s = 1$. We run the test for cases with the numbers of sequences being $n = 40, 100$, respectively. In Figure 5, we plot the probability of error as a function of $\frac{m}{\log n}$. It can be seen that similarly for the test with a reference sequence, as m increases, the probability of error converges to zero. In particular, both curves drop to zero almost at the same threshold, which agrees with Theorem 4.

In Figure 6, we compare the performance of the tests with and without a reference sequence. It can be seen that the probability of error in both curves drops to zero as m increases. More importantly, it is clear that the test without a reference sequence results in an error probability that converges to zero faster than the test with a reference sequence. This is consistent with our theoretical results in Theorems 2 and 5. As we comment in Remark 2, our test without a reference sequence can perform better than the test with a reference

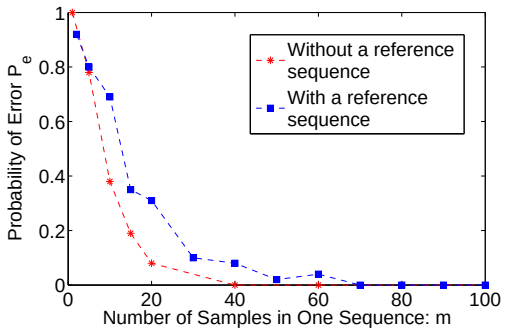

Figure 6: Comparison of the performance of tests with or without a reference sequence.

sequence if α is small enough, where $\alpha = \lim_{n\to\infty} \frac{s}{n}$.

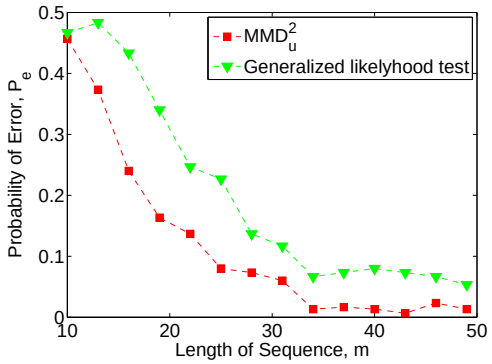

Figure 7: Comparison of the MMD-based test with divergence-based generalized likelihood test with a reference sequence.

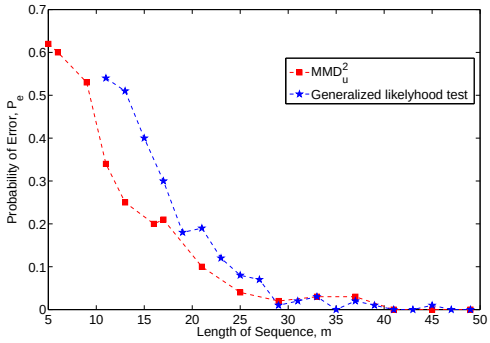

Figure 8: Comparison of the MMD-based test with divergence-based generalized likelihood test without a reference sequence.

5.2 Comparison with Other Tests

In this subsection, we compare our MMD-based tests with tests based on other nonparametric approaches. We first compare our test with the divergence-based generalized likelihood approach developed in [5]. Since the test in [5] is applicable only when the distributions p and q are discrete and have finite alphabets, we set the distributions p and q to be binary with p having probability 0.3 to take "0" (and probability 0.7 to take "1"), and q having probability 0.7 to take "0" (and probability 0.3 to take "1"). We let $s = 1$ and assume that s is known. We let $n = 50$. We compare the performance for both the case with a reference sequence and the case without a reference sequence. We note that for the case with a reference data sequence, we use the empirical distribution of the reference sequence to replace the true distribution p in the divergence-based generalized likelihood test proposed in [5] in order for a fair comparison with the MMD-based test.

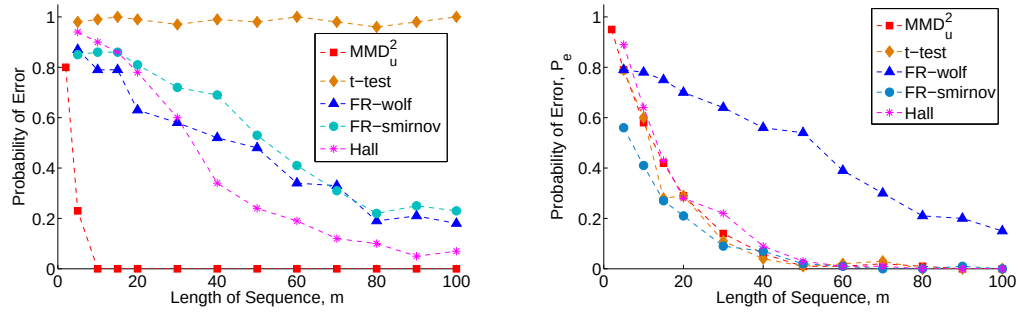

(a) p and q have same mean and same variance

(b) p and q have different mean and same variance

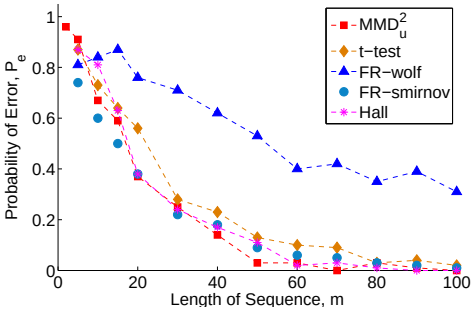

(c) p and q have different mean and different variance

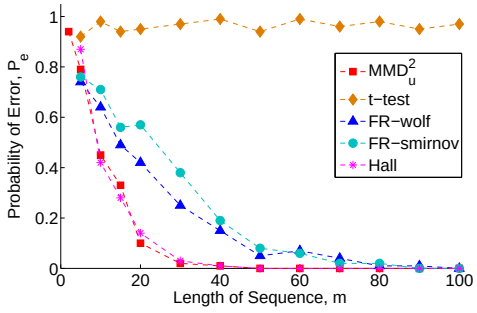

(d) p and q have same mean and different variance

Figure 9: Comparison of the MMD-based test with t-test, FR-Wolf, FR-Smirnov, and Hall tests

In Figure 7 and Figure 8, we plot the probability of error as a function m. It can be seen that the MMD-based test has a slightly better performance than the divergence-based generalized likelihood test in both cases. We note that it has been shown in [5] that the divergence-based test has optimal convergence rate in the limiting case when n is infinite, which suggests that such a test should also perform well for the case with finite n. Thus, the comparison demonstrates that the MMD-based test can provide comparable or even better performance than the well-performed divergence-based test.

We next compare the MMD-based test with four other tests based on traditional statistical approaches which are t-test, FR-Wolf test, FR-Smirnov test, and Hall test. We focus on the scenario with a reference sequence and with $s = 1$. We set $n = 100$. For the MMD-based test, we use the Gaussian kernel with $\sigma = 1$. We compare the five tests for the following four cases:

- p and q have same mean and same variance: p is Laplacian distribution with mean 1 and variance 5, and distribution q is a mixture of two Laplacian distributions with equal probability: one with mean -1 and variance 1, and the other with mean 3 and variance 1.
- p and q have different mean and same variance: p is Gaussian distribution with mean 0 and variance 1, and q is Laplacian distribution with mean 1 and variance 1.
- p and q have different mean and different variance: p is Gaussian distribution with

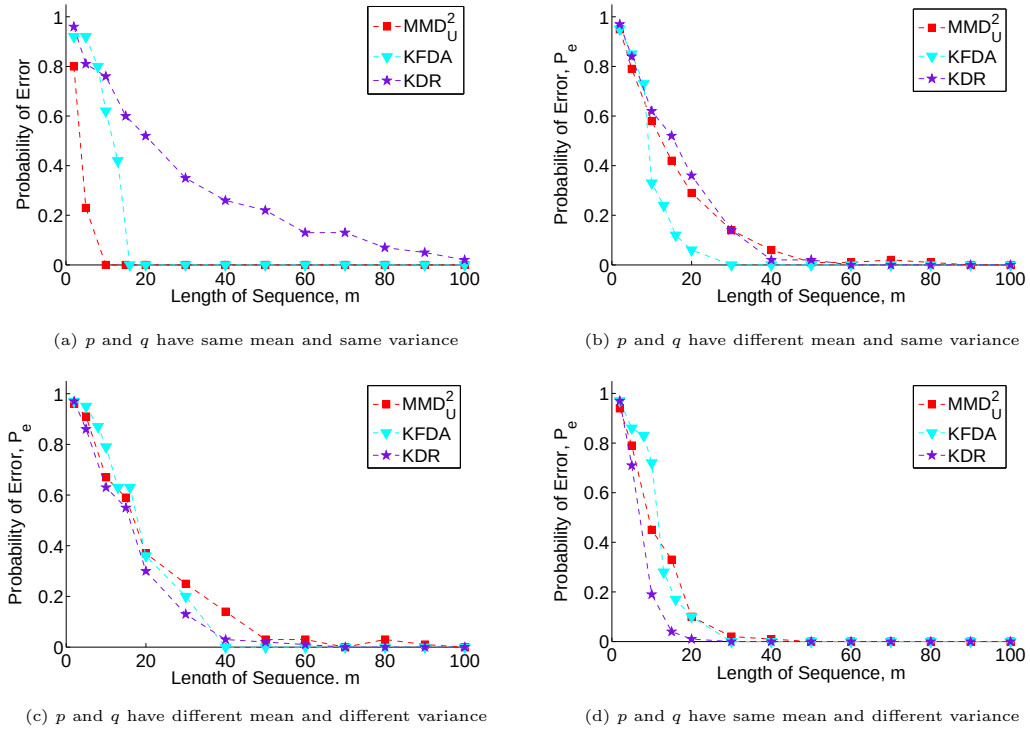

(a) p and q have same mean and same variance

(b) p and q have different mean and same variance

(c) p and q have different mean and different variance

(d) p and q have same mean and different variance

Figure 10: Comparison of the MMD-based test with two kernel-based tests KFDA and KDR

mean 0 and variance 1, and q is Laplacian distribution with mean 1 and variance 3.

- p and q have same mean and different variance: p is Gaussian distribution with mean 0 and variance 1, and q is a mixture of two Laplacian distributions with equal probability: one with mean -2 variance 3 and the other with mean 2 and variance 3.

In Figure 9, we plot the probability of error as a function of m for the five tests for the above four cases. It can be seen that for all cases, the MMD-based test is either the best or one of the best tests among the five tests. In particular, the MMD-based test performs much better than other tests for the case when p and q have the same mean and variance. This suggests that the MMD-based test is specially advantageous to capture difference in higher order moments between two distributions compared to other tests. We also note that although the Hall test sometimes yields comparable performance with the MMD-based test, its complexity is much larger than the MMD-based test.

We next compare our MMD-based tests with two other kernel-based tests KFDA and KDR, where KDR uses the divergence as the distance metric. We focus on the scenario with a reference sequence and with $s = 1$. We set $n = 100$. For all tests, we use the Gaussian kernel with $\sigma = 1$. We compare the three kernel-based tests for the same four cases as in the previous comparison. In Figure 10, we plot the probability of error as a function of m for the three tests for the above four cases. It can be seen that the MMD-based test performs much better than the other two tests when p and q have the same mean and variance, and

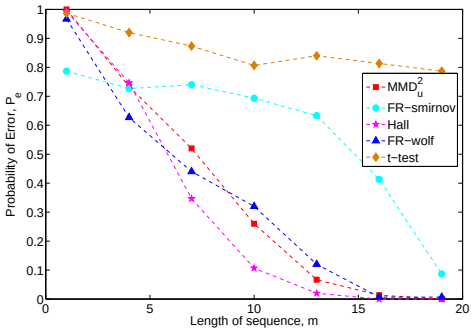

Figure 11: Comparison of the MMD-based test with four other tests on a real data set.

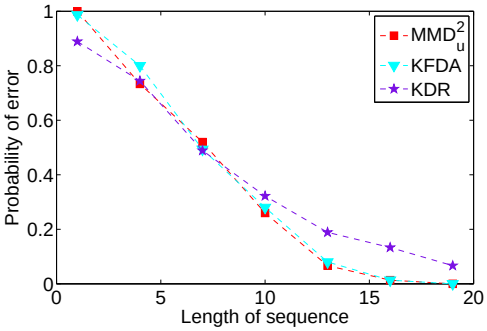

Figure 12: Comparison of the MMD-based test with two other kernel-based tests on a real data set.

performs as well as the other two tests in the remaining three cases.

5.3 Application to Real Data Set

In this subsection, we study how the MMD-based test performs on a real data set. We choose the collection of daily maximum temperature of Syracuse (New York, USA) in July from 1993 to 2012 as the normal data sequences, and the collection of daily maximum temperature of Makapulapai (Hawaii, USA) in May from 1993 to 2012 as anomalous sequences. Here, each data sequence contains daily maximum temperatures of a certain day across twenty years from 1993 to 2012. In our experiment, the data set contains 32 sequences in total, including one temperature sequence of Hawaii and 31 sequences of Syracuse. The probability of error is averaged over all cases with each using one sequence of Hawaii as the anomalous sequence. Although it seems easy to detect the sequence of Hawaii out of the sequences of Syracuse, the temperatures we compare for the two places are in May for Hawaii and July for Syracuse, during which the two places have approximately the same mean in temperature. In this way, it may not be easy to detect the anomalous sequence (in fact, some tests do not perform well as shown in Figure 11).

We apply the MMD-based test (for the case with no reference sequence) to the data set and compare its performance with t-test, FR-Wolf test, FR-Smirnov test, and Hall test. For the MMD-based test, we use the Gaussian kernel with $\sigma = 1$. In Figure 11, we plot the probability of error as a function of the length of sequence m for all tests. It can be seen that the MMD-based test, Hall test, and FR-wolf test have the best performances, and all of the three tests are consistent with the probability of error converging to zero as m goes to infinity. Furthermore, comparing to Hall and FR-wolf tests, the MMD-based test has the lowest computational complexity.

We also compare the performance of MMD-based test with the kernel-based tests KFDA and KDR for the same data set. For all three tests, we use Gaussian kernel with $\sigma = 1$. In Figure 12, we plot the probability of error as a function of the length of sequence for all

tests. It can be seen that all tests are consistent with the probability of error converging to zero as m increases, and the MMD-based test has the best performance among the three tests.

6 Conclusion

In this paper, we have investigated a nonparametric anomaly detection problem, in which normal and anomalous data sequences contain i.i.d. samples drawn from different distributions p and q, respectively. We have built MMD-based distribution-free tests to detect anomalous sequences for two scenarios respectively with and without a reference sequence generated by p. We have characterized the scaling behavior of the sample size m of samples as the total number n of sequences goes to infinity in order to guarantee consistency of the developed tests. We have also compared the scaling behavior of m for the two scenarios and for the cases with and without the knowledge of the number s of anomalous sequences. We have demonstrated our theoretic results by numerical examples and have demonstrated the performance of our tests by comparing it to other appealing tests. Our study of this problem demonstrates a useful application of the mean embedding of distributions and MMD, and we believe that such an approach can be applied to solving various other nonparametric problems.

Appendix

A Proof of Theorem 1

We first introduce the McDiarmid's inequality which is useful in bounding the probability of error in our proof.

Lemma 1 (McDiarmid's Inequality)**.** *Let $f : \mathcal{X}^m \to \mathbb{R}$ be a function such that for all $i \in \{1, \ldots, m\}$, there exist $c_i < \infty$ for which*

$$\sup_{X \in \mathcal{X}^m, \tilde{x} \in \mathcal{X}} |f(x_1, \ldots, x_m) - f(x_1, \ldots x_{i-1}, \tilde{x}, x_{i+1}, \ldots, x_m)| \leq c_i. \tag{23}$$

Then for all probability measure p and every $\epsilon > 0$,

$$P_X\left(f(X) - E_X(f(X)) > \epsilon\right) < \exp\left(-\frac{2\epsilon^2}{\sum_{i=1}^m c_i^2}\right), \tag{24}$$

where X denotes $(x_1, \ldots, x_m)$, E_X denotes the expectation over the m random variables $x_i \sim p$, and P_X denotes the probability over these m variables.

We recall the notation $X := (x_1, x_2, \ldots, x_m)$ and $Y_k := (y_{k1}, y_{k2}, \ldots, y_{km})$. Now for any index $1 \leq a \leq m$, we use X_{-a} and $Y_{k,-a}$ to denote X and Y_k with the a-th component being removed, respectively.

In the test (6), we compute the $\mathrm{MMD}_u^2[X, Y_k]$ for each k, where $1 \leq k \leq n$, as follows:

$$\begin{aligned}\mathrm{MMD}_u^2[X, Y_k] =& \frac{1}{m(m-1)} \sum_{i=1}^{m} \sum_{j \neq i}^{m} k(x_i, x_j) + \frac{1}{m(m-1)} \sum_{i=1}^{m} \sum_{j \neq i}^{m} k(y_{ki}, y_{kj}) \\ & - \frac{2}{m^2} \sum_{i=1}^{m} \sum_{j=1}^{m} k(x_i, y_{kj}). \end{aligned} \tag{25}$$

In order to analyze the error probability for test (6), without loss of generality, we assume that Y_1 is the anomalous sequence, i.e., the sequence generated from the anomalous distribution q. Hence,

$$\begin{aligned} P_e &= P(k^* \neq 1) \\ &= P\bigg(\exists k \neq 1 : \mathrm{MMD}_u^2[X, Y_k] > \mathrm{MMD}_u^2[X, Y_1]\bigg) \\ &\leq \sum_{k=2}^{n} P\bigg(\mathrm{MMD}_u^2[X, Y_k] > \mathrm{MMD}_u^2[X, Y_1]\bigg) \\ &= \sum_{k=2}^{n} P\bigg(\mathrm{MMD}_u^2[X, Y_k] - \mathrm{MMD}_u^2[X, Y_1] + \mathrm{MMD}^2[p, q] > \mathrm{MMD}^2[p, q]\bigg). \end{aligned} \tag{26}$$

For any $k \geq 2$, we define

$$\begin{aligned} \Delta_k =& \mathrm{MMD}_u^2[X, Y_k] - \mathrm{MMD}_u^2[X, Y_1] \\ =& \frac{1}{m(m-1)} \sum_{\substack{i,j=1 \\ i \neq j}}^{m,m} (k(y_{ki}, y_{kj}) - k(y_{1i}, y_{1j})) - \frac{2}{m^2} \sum_{i,j=1}^{m,m} (k(x_i, y_{kj}) - k(x_i, y_{1j})). \end{aligned} \tag{27}$$

It can be seen that $\mathbb{E}[\Delta_k] = -\mathrm{MMD}^2[p, q]$.

In order to apply McDiarmid's inequality to bound (26), we first evaluate the following quantities. There are in total $3m$ variables that affect the value of Δ_k. We study the influence of these $3m$ variables on Δ_k in the following three cases.

For $1 \leq a \leq m$, change of y_{1a} affects Δ_k through the following terms,

$$-\frac{2}{m(m-1)} \sum_{i=1, i \neq a}^{m} k(y_{1i}, y_{1a}) + \frac{2}{m^2} \sum_{i=1}^{m} k(x_i, y_{1a})$$

In this test, the kernel we use is bounded, $0 \leq k(x, y) \leq K$. Hence, we have

$$|\Delta_k(X, Y_k, Y_{1,-a}, y_{1a}) - \Delta_k(X, Y_k, Y_{1,-a}, y'_{1a})| \leq \frac{4K}{m}. \tag{28}$$

Similarly, for $1 \leq a \leq m$, and for $k \geq 2$, change of y_{ka} affects Δ_k through the following terms,

$$\frac{2}{m(m-1)} \sum_{i=1, i \neq a}^{m} k(y_{ki}, y_{ka}) - \frac{2}{m^2} \sum_{i=1}^{m} k(x_i, y_{ka}).$$

Therefore, we have

$$|\Delta_k(X, Y_1, Y_{k,-a}, y_{ka}) - \Delta_k(X, Y_1, Y_{k,-a}, y'_{ka})| \leq \frac{4K}{m}. \tag{29}$$

For $1 \leq a \leq m$, change of x_a affects Δ_k through the following terms,

$$-\frac{2}{m^2}\sum_{j=1}^{m}(k(x_a, y_{kj}) - k(x_a, y_{1j})).$$

Hence, we have

$$|\Delta_k(X_{-a}, x_a, Y_1, Y_k) - \Delta_k(X_{-a}, x'_a, Y_1, Y_k)| \leq \frac{4K}{m}. \tag{30}$$

Combining (28), (29), and (30), and applying McDiarmid's inequality, we have

$$\begin{aligned}
&P\left(\mathrm{MMD}_u^2[X, Y_k] - \mathrm{MMD}_u^2[X, Y_1] + \mathrm{MMD}^2[p, q] > \mathrm{MMD}^2[p, q]\right) \\
&= P\left(\Delta_k - \mathbb{E}[\Delta_k] > \mathrm{MMD}^2[p, q]\right) \\
&\leq \exp\left(-\frac{2\mathrm{MMD}^4[p, q]}{m\frac{16K^2}{m^2} + m\frac{16K^2}{m^2} + m\frac{16K^2}{m^2}}\right) \\
&= \exp\left(-\frac{m\mathrm{MMD}^4[p, q]}{24K^2}\right)
\end{aligned} \tag{31}$$

Therefore, we obtain

$$P_e \leq n\exp\left(-\frac{m\mathrm{MMD}^4[p, q]}{24K^2}\right) = \exp\left(\log n - \frac{m\mathrm{MMD}^4[p, q]}{24K^2}\right), \tag{32}$$

which implies if

$$m > \frac{24K^2(1+\eta)}{\mathrm{MMD}^4[p, q]}\log n,$$

where η is any positive constant, then $P_e \to 0$ as $n \to \infty$. It is also clear that if the above condition is satisfied, P_e converges to zero exponentially fast with respect to m. This completes the proof.

B Proof of Theorem 2

This proof is a generalization of the proof in Appendix A. Without loss of generality, we assume that the first s sequences are generated from the anomalous distribution q. We

analyze the test (8) as follows:

$$\begin{aligned} P_e &= P\bigg(\widehat{\mathcal{I}} \neq \mathcal{I}\bigg) \\ &= P\bigg(\exists k > s : \mathrm{MMD}_u^2[X, Y_k] > \min_{1 \leq l \leq s} \mathrm{MMD}_u^2[X, Y_l]\bigg) \\ &\leq \sum_{k=s+1}^{n} \sum_{l=1}^{s} P\bigg(\mathrm{MMD}_u^2[X, Y_k] > \mathrm{MMD}_u^2[X, Y_l]\bigg) \end{aligned} \tag{33}$$

Following steps similar to those for obtaining (31), we have,

$$P\bigg(\mathrm{MMD}_u^2[X, Y_k] > \mathrm{MMD}_u^2[X, Y_l]\bigg) \leq \exp\bigg(-\frac{m\mathrm{MMD}^4[p, q]}{24K^2}\bigg),$$

for $s+1 \leq k \leq n$ and $1 \leq l \leq s$. Hence,

$$P_e \leq \exp\left(\log((n-s)s) - \frac{m\mathrm{MMD}^4[p, q]}{24K^2}\right), \tag{34}$$

which implies that if

$$m > \frac{24K^2(1+\eta)}{\mathrm{MMD}^4[p, q]} \log((n-s)s),$$

where η is any positive constant, then $P_e \to 0$ as $n \to \infty$.

It is also clear that if the above condition is satisfied, P_e converges to zero exponentially fast with respect to m.

The computations of the test include $\mathcal{O}(nm^2)$ computations of the kernel function values $k(x, y)$, $\mathcal{O}(nm^2)$ additions, $\mathcal{O}(n)$ multiplications, and $\mathcal{O}(ns - \frac{s^2}{2})$ comparisons. Hence, the computational complexity is $\max\{\mathcal{O}(nm^2), \mathcal{O}(ns - \frac{s^2}{2})\}$.

C Proof of Theorem 3

We analyze the probability of error for the test (10) as follows. When $s \neq 0$, without loss of generality, we assume that the first s sequences are the anomalous sequences. Hence,

$$\begin{aligned} P_e &= P\bigg(\widehat{\mathcal{I}} \neq \mathcal{I}\bigg) \\ &= P\bigg(\big(\exists 1 \leq l \leq s : \mathrm{MMD}_u^2[X, Y_l] \leq \delta_n\big) or \big(\exists s+1 \leq k \leq n : \mathrm{MMD}_u^2[X, Y_k] > \delta_n\big)\bigg) \\ &\leq \sum_{l=1}^{s} P\bigg(\mathrm{MMD}_u^2[X, Y_l] \leq \delta_n\bigg) + \sum_{k=s+1}^{n} P\bigg(\mathrm{MMD}_u^2[X, Y_k] > \delta_n\bigg). \end{aligned} \tag{35}$$

Since $\delta_n \to 0$ as $n \to \infty$, for large enough n, $\mathrm{MMD}^2[p,q] - \delta_n > 0$ holds. It can also be seen that $\mathbb{E}\big[\mathrm{MMD}_u^2[X,Y_l]\big] = \mathrm{MMD}^2[p,q]$, and $\mathbb{E}\big[\mathrm{MMD}_u^2[X,Y_k]\big] = 0$. Hence, applying McDiarmid's inequality, we obtain,

$$\begin{aligned} P\bigg(\mathrm{MMD}_u^2[X,Y_l] \leq \delta_n\bigg) &= P\bigg(\mathrm{MMD}_u^2[X,Y_l] - \mathrm{MMD}^2[p,q] \leq -(\mathrm{MMD}^2[p,q] - \delta_n)\bigg) \\ &\leq \exp\bigg(-\frac{2(\mathrm{MMD}^2[p,q]-\delta_n)^2}{\frac{16K^2}{m}+\frac{16K^2}{m}}\bigg) \\ &= \exp\bigg(-\frac{m(\mathrm{MMD}^2[p,q]-\delta_n)^2}{16K^2}\bigg), \end{aligned} \tag{36}$$

for $1 \leq l \leq s$, and,

$$P\bigg(\mathrm{MMD}_u^2[X,Y_k] > \delta_n\bigg) \leq \exp\bigg(-\frac{2\delta_n^2}{\frac{16K^2}{m}+\frac{16K^2}{m}}\bigg) = \exp\bigg(-\frac{m\delta_n^2}{16K^2}\bigg), \tag{37}$$

for $s+1 \leq k \leq n$. Therefore,

$$\begin{aligned} P_e \leq & s\exp\bigg(-\frac{m(\mathrm{MMD}^2[p,q]-\delta_n)^2}{16K^2}\bigg) + (n-s)\exp\bigg(-\frac{m\delta_n^2}{16K^2}\bigg) \\ = & \exp\bigg(\log s - \frac{m(\mathrm{MMD}^2[p,q]-\delta_n)^2}{16K^2}\bigg) + \exp\bigg(\log(n-s) - \frac{m\delta_n^2}{16K^2}\bigg). \end{aligned} \tag{38}$$

Hence, if

$$m > \frac{16(1+\eta)K^2}{(\mathrm{MMD}^2[p,q]-\delta_n)^2}\log s, \tag{39}$$

and

$$m > \frac{16(1+\eta)K^2}{\delta_n^2}\log(n-s), \tag{40}$$

where η is any positive constant, then $P_e \to 0$, as $n \to \infty$.

When $s = 0$, $P_e = \sum_{i=1}^{n} P\left(\mathrm{MMD}_u^2[X,Y_k] > \delta_n\right)$. Then applying (37), we have if

$$m > \frac{16(1+\eta)K^2}{\delta_n^2}\log n, \tag{41}$$

where η is any positive constant, then $P_e \to 0$, as $n \to \infty$.

The computations of the test include $\mathcal{O}(nm^2)$ computations of the kernel function values $k(x,y)$, $\mathcal{O}(nm^2)$ additions, $\mathcal{O}(n)$ multiplications, and $\mathcal{O}(n)$ comparisons. Hence, the computational complexity is $\mathcal{O}(nm^2)$.

D Proof of Theorem 4

In order to analyze the probability of error for the test (15), without loss of generality, we assume that the first sequence is the anomalous sequence generated by the anomalous distribution q. Hence,

$$\begin{aligned} P_e &= P(k^* \neq 1) \\ &= P\bigg(\exists k \neq 1 : \mathrm{MMD}_u^2[Y_k, \overline{Y}_k] > \mathrm{MMD}_u^2[Y_1, \overline{Y}_1]\bigg) \\ &\leq \sum_{k=2}^{n} P\bigg(\mathrm{MMD}_u^2[Y_k, \overline{Y}_k] > \mathrm{MMD}_u^2[Y_1, \overline{Y}_1]\bigg), \end{aligned} \tag{42}$$

For notational convenience, we stack $Y_1, \ldots, Y_n$ into a nm dimensional row vector $Y = \{y_i, 1 \leq i \leq nm\}$, where $Y_k = \{y_{(k-1)m+1}, \ldots, y_{km}\}$. And we define $n' = (n-1)m$. We then have,

$$\mathrm{MMD}_u^2[Y_1, \overline{Y}_1] = \frac{1}{m(m-1)} \sum_{\substack{i,j=1 \\ i \neq j}}^{m,m} k(y_i, y_j) + \frac{1}{n'(n'-1)} \sum_{\substack{i,j=m+1 \\ i \neq j}}^{nm} k(y_i, y_j) - \frac{2}{mn'} \sum_{\substack{i=1 \\ j=m+1}}^{m,nm} k(y_i, y_j). \tag{43}$$

For $2 \leq k \leq n$, we have,

$$\begin{aligned} \mathrm{MMD}_u^2[Y_k, \overline{Y}_k] = \frac{1}{m(m-1)} &\sum_{\substack{i,j=(k-1)m+1 \\ i \neq j}}^{km,km} k(y_i, y_j) + \frac{1}{n'(n'-1)} \bigg(\sum_{\substack{i,j=1 \\ i \neq j}}^{m,m} k(y_i, y_j) + 2 \sum_{\substack{i=1 \\ j=m+1}}^{m,(k-1)m} k(y_i, y_j) \\ &+ 2 \sum_{\substack{i=1 \\ j=km+1}}^{m,nm} k(y_i, y_j) + \sum_{\substack{i,j=m+1 \\ i \neq j}}^{(k-1)m,(k-1)m} k(y_i, y_j) + \sum_{\substack{i,j=km+1 \\ i \neq j}}^{nm,nm} k(y_i, y_j) + 2 \sum_{\substack{i=m+1 \\ j=km+1}}^{(k-1)m,nm} k(y_i, y_j) \bigg) \\ &- \frac{2}{mn'} \bigg(\sum_{\substack{i=1 \\ j=(k-1)m+1}}^{m,km} k(y_i, y_j) + \sum_{\substack{i=m+1 \\ j=(k-1)m+1}}^{(k-1)m,km} k(y_i, y_j) + \sum_{\substack{i=(k-1)m+1 \\ j=km+1}}^{km,nm} k(y_i, y_j) \bigg). \end{aligned} \tag{44}$$

We define

$$\Delta_k = \mathrm{MMD}_u^2[Y_k, \overline{Y}_k] - \mathrm{MMD}_u^2[Y_1, \overline{Y}_1].$$

It can be shown that,

$$\mathbb{E}[\mathrm{MMD}_u^2[Y_1, \overline{Y}_1]] = \mathrm{MMD}^2[p, q],$$

and

$$\begin{aligned}\mathbb{E}[\mathrm{MMD}_u^2[Y_k,\overline{Y}_k]] &= \mathbb{E}_{x,x'}k(x,x') + \frac{1}{(n-1)m((n-1)m-1)}\bigg(m(m-1)\mathbb{E}_{y,y'}k(y,y') \\ &\quad + 2m^2(n-2)\mathbb{E}_{x,y}k(x,y) + ((n-2)m-1)(n-2)m\mathbb{E}_{x,x'}k(x,x')\bigg) \\ &\quad - \frac{2}{(n-1)m^2}\bigg(m^2\mathbb{E}_{x,y}k(x,y) + (n-2)m^2\mathbb{E}_{x,x'}k(x,x')\bigg) \\ &\quad \to 0, \text{ as } n\to\infty, \end{aligned} \tag{45}$$

where x and x' are independent but have the same distribution p, y and y' are independent but have the same distribution q. Hence, there exists a constant θ that satisfies

$$\mathbb{E}[\mathrm{MMD}_u^2[Y_k,\overline{Y}_k]] < \theta < \mathrm{MMD}^2[p,q], \tag{46}$$

for large enough n.

We next divide the entries in $\{y_1,\ldots,y_{nm}\}$ into three groups: $Y_1 = \{y_1,\ldots,y_m\}$, $Y_k = \{y_{(k-1)m+1}\ldots,y_{km}\}$, and $\widehat{Y}_k$ that contains the remaining entries. We define Y_{-a} as Y with the a-th component y_a being removed.

For $1\le a\le m$, y_a affects Δ_k through the following terms

$$\begin{aligned}&\frac{1}{n'(n'-1)}\bigg(2\sum_{\substack{j=1\\j\neq a}}^{m}k(y_a,y_j) + 2\sum_{j=m+1}^{(k-1)m}k(y_a,y_j) + 2\sum_{j=km+1}^{nm}k(y_a,y_j)\bigg) \\ &- \frac{2}{mn'}\sum_{j=(k-1)m+1}^{km}k(y_a,y_j) - \frac{2}{m(m-1)}\sum_{\substack{j=1\\k\neq a}}^{m}k(y_a,y_j) + \frac{2}{mn'}\sum_{j=m+1}^{nm}k(y_a,y_j). \end{aligned} \tag{47}$$

Hence, for $1\le a\le m$, we have

$$|\Delta_k\big(Y_{-a},y_a\big) - \Delta_k\big(Y_{-a},y_a'\big)| \le \frac{4K}{m}\bigg(1+\mathcal{O}\bigg(\frac{1}{n}\bigg)\bigg). \tag{48}$$

For $(k-1)m+1\le a\le km$, y_a affects Δ_k through

$$\begin{aligned}&\frac{2}{m(m-1)}\sum_{\substack{j=(k-1)m+1\\j\neq a}}^{km}k(y_a,y_j) - \frac{2}{mn'}\bigg(\sum_{i=1}^{m}k(y_i,y_a) + \sum_{i=m+1}^{(k-1)m}k(y_i,y_a) + \sum_{j=km+1}^{nm}k(y_a,y_j)\bigg) \\ &- \frac{2}{n'(n'-1)}\sum_{\substack{j=m+1\\j\neq a}}^{nm}k(y_a,y_j) + \frac{2}{mn'}\sum_{i=1}^{m}k(y_a,y_i). \end{aligned} \tag{49}$$

Hence, for $(k-1)m+1\le a\le km$, we have

$$|\Delta_k\big(Y_{-a},y_a\big) - \Delta_k\big(Y_{-a},y_a'\big)| \le \frac{4K}{m}\bigg(1+\mathcal{O}\bigg(\frac{1}{n}\bigg)\bigg). \tag{50}$$

For $m+1 \leq a \leq (k-1)m$ or $km+1 \leq a \leq nm$, y_a affects Δ_k through

$$\frac{2}{n'(n'-1)}\Bigg(\sum_{i=1}^{m} k(y_i, y_a) + \sum_{\substack{i=m+1 \\ i\neq a}}^{(k-1)m} k(y_i, y_a) + \sum_{j=km+1}^{nm} k(y_a, y_j)\Bigg) - \frac{2}{mn'} \sum_{j=(k-1)m+1}^{km} k(y_a, y_j)$$

$$- \frac{1}{n'(n'-1)} \sum_{\substack{j=m+1 \\ j\neq a}}^{nm} k(y_a, y_j) + \frac{2}{mn'} \sum_{i=(k-1)m+1}^{km} k(y_i, y_a) \tag{51}$$

Hence, for $m+1 \leq a \leq (k-1)m$ or $km+1 \leq a \leq nm$, we have

$$|\Delta_k\big(Y_{-a}, y_a\big) - \Delta_k\big(Y_{-a}, y'_a\big)| \leq \frac{1}{m}\left(\mathcal{O}\left(\frac{1}{n}\right)\right). \tag{52}$$

We further derive the following probability,

$$\begin{aligned}
&P\Big(\mathrm{MMD}_u^2[Y_k, \overline{Y}_k] > \mathrm{MMD}_u^2[Y_1, \overline{Y}_1]\Big) \\
&= P\Big(\mathrm{MMD}_u^2[Y_k, \overline{Y}_k] - \mathrm{MMD}_u^2[Y_1, \overline{Y}_1] + \mathrm{MMD}^2[p,q] > \mathrm{MMD}^2[p,q]\Big) \\
&\overset{(a)}{\leq} P\Big(\mathrm{MMD}_u^2[Y_k, \overline{Y}_k] - \mathrm{MMD}_u^2[Y_1, \overline{Y}_1] + \mathrm{MMD}^2[p,q] - \mathbb{E}[\mathrm{MMD}_u^2[Y_k, \overline{Y}_k]] > \mathrm{MMD}^2[p,q] - \theta\Big),
\end{aligned} \tag{53}$$

where (a) follows from (46).

Combining (48), (50), (52), and applying McDiarmid's inequality, we have,

$$\begin{aligned}
&P\Big(\mathrm{MMD}_u^2[Y_k, \overline{Y}_k] > \mathrm{MMD}_u^2[Y_1, \overline{Y}_1]\Big) \\
&\qquad \leq \exp\left(-\frac{2(\mathrm{MMD}^2[p,q]-\theta)^2}{2m\frac{16K^2}{m^2}(1+\mathcal{O}(\frac{1}{n})) + \frac{1}{m}\mathcal{O}(\frac{1}{n})}\right) \\
&\qquad = \exp\left(-\frac{m(\mathrm{MMD}^2[p,q]-\theta)^2}{16K^2(1+\mathcal{O}(\frac{1}{n}))}\right)
\end{aligned} \tag{54}$$

Hence,

$$P_e \leq \exp\left(\log n - \frac{m(\mathrm{MMD}^2[p,q]-\theta)^2}{16K^2(1+\mathcal{O}(\frac{1}{n}))}\right). \tag{55}$$

Since θ can be picked arbitrarily close to zero, we conclude that if

$$m > \frac{16K^2(1+\eta)}{\mathrm{MMD}^4[p,q]} \log n, \tag{56}$$

where η is any positive constant, then $P_e \to 0$ as $n \to \infty$. It is also clear that if the above condition is satisfied, P_e converges to zero exponentially fast with respect to m. This completes the proof.

E Proof of Theorem 5

We analyze the performance of the test (17). Without loss of generality, we assume that the first s sequences are anomalous and are generated from distribution q. Hence, the probability of error can be bounded as,

$$\begin{aligned} P_e =& P\bigg(\exists k > s : \mathrm{MMD}_u^2[Y_k, \overline{Y}_k] \quad > \min_{1\le l\le s} \mathrm{MMD}_u^2[Y_l, \overline{Y}_l]\bigg) \\ \le& \sum_{k=s+1}^{n}\sum_{l=1}^{s} P\bigg(\mathrm{MMD}_u^2[Y_k, \overline{Y}_k] > \mathrm{MMD}_u^2[Y_l, \overline{Y}_l]\bigg). \end{aligned} \tag{57}$$

Using the fact that $\frac{s}{n} \to \alpha$, where $0 \le \alpha < \frac{1}{2}$, and using (43) and (44), we can show that

$$\mathbb{E}\big[\mathrm{MMD}_u^2[Y_l, \overline{Y}_l]\big] \to (1-\alpha)^2\mathrm{MMD}^2[p,q], \tag{58}$$

as $n \to \infty$, for $1 \le l \le s$, and

$$\mathbb{E}\big[\mathrm{MMD}_u^2[Y_k, \overline{Y}_k]\big] \to \alpha^2\mathrm{MMD}^2[p,q], \tag{59}$$

as $n \to \infty$, for $s+1 \le k \le n$. Hence, there exists a constant θ such that

$$0 < \theta < (1-\alpha)^2\mathrm{MMD}^2[p,q] - \alpha^2\mathrm{MMD}^2[p,q]$$

and

$$\mathbb{E}\big[\mathrm{MMD}_u^2[Y_k, \overline{Y}_k] - \mathrm{MMD}_u^2[Y_l, \overline{Y}_l]\big] < \alpha^2\mathrm{MMD}^2[p,q] - (1-\alpha)^2\mathrm{MMD}^2[p,q] + \theta, \tag{60}$$

for large enough n.

Therefore, we obtain,

$$\begin{aligned} & P\bigg(\mathrm{MMD}_u^2[Y_k, \overline{Y}_k] - \mathrm{MMD}_u^2[Y_l, \overline{Y}_l] > 0\bigg) \\ =& P\bigg(\mathrm{MMD}_u^2[Y_k, \overline{Y}_k] - \mathrm{MMD}_u^2[Y_l, \overline{Y}_l] - \mathbb{E}\big[\mathrm{MMD}_u^2[Y_k, \overline{Y}_k] - \mathrm{MMD}_u^2[Y_l, \overline{Y}_l]\big] \\ & \qquad > -\mathbb{E}\big[\mathrm{MMD}_u^2[Y_k, \overline{Y}_k] - \mathrm{MMD}_u^2[Y_l, \overline{Y}_l]\big]\bigg) \\ \le& P\bigg(\mathrm{MMD}_u^2[Y_k, \overline{Y}_k] - \mathrm{MMD}_u^2[Y_l, \overline{Y}_l] - \mathbb{E}\big[\mathrm{MMD}_u^2[Y_k, \overline{Y}_k] - \mathrm{MMD}_u^2[Y_l, \overline{Y}_l]\big] \\ & \qquad > ((1-\alpha)^2 - \alpha^2)\mathrm{MMD}^2[p,q]) - \theta\bigg), \end{aligned} \tag{61}$$

for large enough n.

Applying McDiarmid's inequality, we obtain,

$$P_e \le \exp\left(\log((n-s)s) - \frac{m((1-2\alpha)\mathrm{MMD}^2[p,q] - \theta)^2}{16K^2(1+\mathcal{O}(\frac{1}{n}))}\right). \tag{62}$$

Since θ can be arbitrarily small, we conclude that if,

$$m > \frac{16K^2(1+\eta)}{(1-2\alpha)^2\mathrm{MMD}^4[p,q]}\log(s(n-s)), \tag{63}$$

where η is any positive constant, then $P_e \to 0$, as $n \to \infty$. It is also clear that if the above condition is satisfied, P_e converges to zero exponentially fast with respect to m.

The computations of the test (17) include $\mathcal{O}(n^3m^2)$ computations of kernel function values $k(x,y)$, $\mathcal{O}(n^3m^2)$ additions, $\mathcal{O}(n)$ multiplications, and $\mathcal{O}(ns - \frac{s^2}{2})$ comparisons. Hence, the computational complexity is $\mathcal{O}(n^3m^2)$.

F Proof of Theorem 6

We analyze the performance of the test (19). Without loss of generality, we assume that the first s sequences are the anomalous sequences. Hence,

$$\begin{aligned} P_e &= P\bigg(\big(\exists 1 \le l \le s : \mathrm{MMD}_u^2[Y_l, \overline{Y}_l] \le \delta_n\big) \text{or } \big(\exists s+1 \le k \le n : \mathrm{MMD}_u^2[Y_k, \overline{Y}_k] > \delta_n\big)\bigg) \\ &\le \sum_{l=1}^{s} P\bigg(\mathrm{MMD}_u^2[Y_l, \overline{Y}_l] \le \delta_n\bigg) + \sum_{k=s+1}^{n} P\bigg(\mathrm{MMD}_u^2[Y_k, \overline{Y}_k] > \delta_n\bigg). \end{aligned} \tag{64}$$

Using the fact that $\frac{s}{n} \to 0$ as $n \to \infty$, and using (43) and (44) we obtain,

$$\mathbb{E}\big[\mathrm{MMD}_u^2[Y_l, \overline{Y}_l]\big] \to \mathrm{MMD}^2[p,q], \tag{65}$$

$$\mathbb{E}\big[\mathrm{MMD}_u^2[Y_k, \overline{Y}_k]\big] \to 0, \tag{66}$$

as $n \to \infty$, for $1 \le l \le s$ and $s+1 \le k \le n$.

Due to (65), there exists a constant ϵ such that $-\mathbb{E}\big[\mathrm{MMD}_u^2[Y_l, \overline{Y}_l]\big] < -\mathrm{MMD}^2[p,q] + \epsilon$ for large enough n.

For $1 \le l \le s$, we drive,

$$\begin{aligned} &P\bigg(\mathrm{MMD}_u^2[Y_l, \overline{Y}_l] \le \delta_n\bigg) \\ &= P\bigg(\mathrm{MMD}_u^2[Y_l, \overline{Y}_l] - \mathbb{E}\big[\mathrm{MMD}_u^2[Y_l, \overline{Y}_l]\big] \le -\mathbb{E}\big[\mathrm{MMD}_u^2[Y_l, \overline{Y}_l]\big] + \delta_n\bigg) \\ &\le P\bigg(\mathrm{MMD}_u^2[Y_l, \overline{Y}_l] - \mathbb{E}\big[\mathrm{MMD}_u^2[Y_l, \overline{Y}_l] \le -(\mathrm{MMD}^2[p,q] - \epsilon - \delta_n)\bigg), \end{aligned} \tag{67}$$

for large enough n. Therefore, by applying McDiarmid's inequality, we obtain,

$$\begin{aligned}
&P\Big(\mathrm{MMD}_u^2[Y_l, \overline{Y}_l] \leq \delta_n\Big) \\
&\quad \leq \exp\bigg(-\frac{2(\mathrm{MMD}^2[p,q] - \epsilon - \delta_n)^2}{\frac{16K^2}{m}(1+\mathcal{O}(\frac{1}{n})) + \frac{16K^2}{m}(1+\mathcal{O}(\frac{1}{n}))}\bigg) \\
&\quad = \exp\bigg(-\frac{m(\mathrm{MMD}^2[p,q] - \epsilon - \delta_n)^2}{16K^2(1+\mathcal{O}(\frac{1}{n}))}\bigg), \qquad (68)
\end{aligned}$$

for large enough n.

For $s+1 \leq k \leq n$,

$$\begin{aligned}
&P\Big(\mathrm{MMD}_u^2[Y_k, \overline{Y}_k] > \delta_n\Big) \\
&\quad = P\Big(\mathrm{MMD}_u^2[Y_k, \overline{Y}_k] - \mathbb{E}\big[\mathrm{MMD}_u^2[Y_k, \overline{Y}_k]\big] > \delta_n - \mathbb{E}\big[\mathrm{MMD}_u^2[Y_k, \overline{Y}_k]\big]\Big). \qquad (69)
\end{aligned}$$

Using the fact that $\frac{s^2}{n^2\delta_n} \to 0$ as $n \to \infty$, we can show that

$$\frac{\mathbb{E}\big[\mathrm{MMD}_u^2[Y_k, \overline{Y}_k]\big]}{\delta_n} \to 0,$$

as $n \to \infty$. Hence, for large enough n, $\delta_n > \mathbb{E}\big[\mathrm{MMD}_u^2[Y_k, \overline{Y}_k]\big]$. Therefore, using McDiarmid's inequality, we have

$$\begin{aligned}
&P\Big(\mathrm{MMD}_u^2[Y_k, \overline{Y}_k] > \delta_n\Big) \\
&\quad \leq \exp\bigg(-\frac{2(\delta_n - \mathbb{E}\big[\mathrm{MMD}_u^2[Y_k, \overline{Y}_k]\big])^2}{\frac{16K^2}{m}(1+\mathcal{O}(\frac{1}{n})) + \frac{16K^2}{m}(1+\mathcal{O}(\frac{1}{n}))}\bigg) \\
&\quad = \exp\bigg(-\frac{m(\delta_n - \mathbb{E}\big[\mathrm{MMD}_u^2[Y_k, \overline{Y}_k]\big])^2}{16K^2(1+\mathcal{O}(\frac{1}{n}))}\bigg). \qquad (70)
\end{aligned}$$

Therefore,

$$\begin{aligned}
P_e &\leq s\exp\bigg(-\frac{m(\mathrm{MMD}^2[p,q] - \epsilon - \delta_n)^2}{16K^2(1+\mathcal{O}(\frac{1}{n}))}\bigg) \\
&\quad + (n-s)\exp\bigg(-\frac{m(\delta_n - \mathbb{E}\big[\mathrm{MMD}_u^2[Y_k, \overline{Y}_k]\big])^2}{16K^2(1+\mathcal{O}(\frac{1}{n}))}\bigg) \\
&= \exp\bigg(\log s - \frac{m(\mathrm{MMD}^2[p,q] - \epsilon - \delta_n)^2}{16K^2(1+\mathcal{O}(\frac{1}{n}))}\bigg) \\
&\quad + \exp\bigg(\log(n-s) - \frac{m(\delta_n - \mathbb{E}\big[\mathrm{MMD}_u^2[Y_k, \overline{Y}_k]\big])^2}{16K^2(1+\mathcal{O}(\frac{1}{n}))}\bigg), \qquad (71)
\end{aligned}$$

for large enough n. Hence, we conclude that if

$$m > \frac{16(1+\eta)K^2}{(\mathrm{MMD}^2[p,q] - \delta_n)^2} \log s, \tag{72}$$

and

$$m > \frac{16(1+\eta)K^2}{(\delta_n - \mathbb{E}\big[\mathrm{MMD}_u^2[Y_k, \overline{Y}_k]\big])^2} \log(n-s), \tag{73}$$

where η is any positive constant, then $P_e \to 0$, as $n \to \infty$.

When $s = 0$, $P_e = \sum_{k=1}^{n} P(\mathrm{MMD}_u^2[Y_k, \overline{Y_k}] > \delta_n)$, then applying (70), we have if

$$m > \frac{16(1+\eta)K^2}{(\delta_n - \mathbb{E}\big[\mathrm{MMD}_u^2[Y_k, \overline{Y}_k]\big])^2} \log n, \tag{74}$$

where η is any positive constant, then $P_e \to 0$, as $n \to \infty$.

The computations of the test (19) include $\mathcal{O}(n^3m^2)$ computations of kernel function values $k(x,y)$, $\mathcal{O}(n^3m^2)$ additions, $\mathcal{O}(n)$ multiplications, and $\mathcal{O}(n)$ comparisons. Hence, the computational complexity is $\mathcal{O}(n^3m^2)$.

References

[1] L. Lai, H. V. Poor, Y. Xin, and G. Georgiadis. Quickest search over multiple sequences. *IEEE Trans. Inform. Theory*, 57(8):5375–5386, August 2011.

[2] A. Tajer and H. V. Poor. Quick search for rare events. *IEEE Trans. Inform. Theory*, 59(7):4462–4481, July 2013.

[3] A. Patcha and J.-M. Park. An overview of anomaly detection techniques: Existing solutions and latest technological trends. *Comput. Networks*, 51(12):3448–3470, August 2007.

[4] V. Chandola, A. Banerjee, and V. Kumar. Anomaly detection: A survey. *ACM Computing Surveys*, 41(3):1–58, July 2009.

[5] Y. Li, S. Nitinawarat, and V.V. Veeravalli. Universal outlier hypothesis testing. Submitted to *IEEE Trans. Inform. Theory, May 2013*.

[6] J. H. Friedman and L. C. Rafsky. Multivariate generalizations of the wald-wolfowitz and smirnov two-sample tests. *Ann. Statist.*, 7(4):pp. 697–717, 1979.

[7] P. Hall and N. Tajvidi. Permutation tests for equality of distributions in high-dimensional settings. *Biometrika*, 89(2):pp. 359–374, 2002.

[8] T. Kanamori, T. Suzuki, and M. Sugiyama. Divergence estimation and two-sample homogeneity test under semiparametric density-ratio models. *IEEE Trans. Inform. Theory*, 58(2):708–720, Feb 2012.

[9] Z. Harchaoui, F. Bach, and E. Moulines. Testing for homogeneity with kernel fisher discriminant analysis. In *Proc. Advances in Neural Information Processing Systems (NIPS)*, 2008.

[10] A. Berlinet and C. Thomas-Agnan. *Reproducing Kernel Hilbert Spaces in Probability and Statistics*. Kluwer, 2004.

[11] B. Sriperumbudur, A. Gretton, K. Fukumizu, G. Lanckriet, and B. Schölkopf. Hilbert space embeddings and metrics on probability measures. *J. Mach. Learn. Res.*, 11:1517–1561, 2010.

[12] K. Fukumizu, A. Gretton, X. Sun, and B. Schölkopf. Kernel measures of conditional dependence. In *Proc. Advances in Neural Information Processing Systems (NIPS)*, 2008.

[13] B. Sriperumbudur, A. Gretton, K. Fukumizu, G. Lanckriet, and B. Schölkopf. Injective Hilbert space embeddings of probability measures. In *Proc. Annual Conference on Learning Theory (COLT)*, 2008.

[14] K. Fukumizu, B. Sriperumbudur, A. Gretton, and B. Schölkopf. Characteristic kernels on groups and semigroups. In *Proc. Advances in Neural Information Processing Systems (NIPS)*, 2009.

[15] A. Gretton, K. Borgwardt, M. Rasch, B. Schökopf, and A. Smola. A kernel method for the two-sample-problem. In *Proc. Advances in Neural Information Processing Systems (NIPS)*, 2007.

[16] A. Gretton, K. Borgwardt, M. Rasch, B. Schölkopf, and A. Smola. A kernel two-sample test. *J. Mach. Learn. Res.*, 13:723–773, 2012.

[17] B. Schölkopf and A. J. Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. The MIT Press, Cambridge, MA, USA, 2001.

[18] T. Hofmann, B. Schölkopf, and A. J. Smola. Kernel methods in machine learning. *Ann. Statist.*, 36(3):1171–1220, 2008.